

The Large Scale Curvature of Networks

Onuttom Narayan¹ and Iraj Saniee²

¹ *Department of Physics, University of California, Santa Cruz, CA 95064 and*

² *Mathematics of Networks Department, Bell Laboratories,
Alcatel-Lucent, 600 Mountain Avenue, Murray Hill, NJ 07974*

(Dated: July 9, 2009)

Understanding key structural properties of large scale networks are crucial for analyzing and optimizing their performance, and improving their reliability and security. Here we show that these networks possess a previously unnoticed feature, global curvature, which we argue has a major impact on core congestion: the load at the core of a network with N nodes scales as N^2 as compared to $N^{1.5}$ for a flat network. We substantiate this claim through analysis of a collection of real data networks across the globe as measured and documented by previous researchers.

PACS numbers:

Large-scale data networks form the infrastructure for contemporary global communications. Increasingly, the trend in these networks is towards converged services over the Internet protocol, dynamic and automatic reconfigurability, and flatter architecture for fast service creation and survivability. In such a large and fast changing environment, there is a need for identifying key structural properties that affect their performance, reliability and security and which provide efficient and scalable models to estimate these metrics reliably.

Recent models of networks have focused on features such as their ‘small world’ property [1, 2, 3] or power law degree distributions [4, 5, 6]. There has been evidence for power-law degree distributions in data networks at the IP layer [7], for the worldwide web [4], and even for the virtual network of social connections [8], but are found not to exist for physical networks such as electrical grids [3, 9] and some biological networks [9, 10]. Although these features are interesting and important, the impact of intrinsic geometrical and topological features of large-scale networks on performance, reliability and security is of much greater importance. Intuitively, it is known that traffic between nodes tends to go through a relatively small core of the network [11] as if the shortest path between them is curved inwards. It has been suggested that this property may be due to *global curvature* of the network [12].

In this paper, we define the global (negative) curvature for finite networks and demonstrate its existence at the IP layer by examining topologies of numerous publicly available networks [13]. A recent report [14], also refers to curvature as a possible cause of some key observations about networks at the IP layer. However, these authors *assume* negative curvature, and construct a model with a few extra simple assumptions that shares various features with real networks such as a power law degree distribution. By contrast, we *demonstrate* negative curvature through direct measurement.

Turning to the impact of negative curvature, we focus on the load (also referred to as the betweenness centrality), as defined by assuming unit traffic between each pair of nodes in the network with shortest path routing, and

calculating the traffic through each node. (This is *not* the actual time-variable demand that is routed through nodes and links at the IP layer.) We show that network curvature or δ -hyperbolicity [15] implies that the load at the core of the network scales with the number of nodes N as $\sim N^2$, which is faster than the $\sim N^{1.5}$ scaling for flat networks. Thus core congestion is worse in hyperbolic networks, and geodesic routing achieved with greedy algorithms on hyperbolic networks [14] is actually problematic. Previous work [16, 17, 18] has considered the load as a function of node degree for fixed N , which we have also examined separately [19].




FIG. 1: A rendering of the graph for the network 7018(AT&T).

Negative curvature of a geodesic metric space is defined by Gromov [15] in terms of the ‘ δ -Thin Triangle Condition’. For a graph, an appropriate metric can be used. For any three nodes (ijk) , the geodesics g_{ij}, g_{jk} and g_{ki} of lengths d_{ij}, d_{jk} and d_{ki} are constructed. A fourth node m is chosen, and the shortest distance between m and all the nodes on (ij) is defined as $d(m; ij)$. The distance $D(m; ijk)$ is defined as the maximum of

$d(m; ij), d(m; jk)$ and $d(m; ki)$. Then if

$$\max_{(ijk)} \min_m D(m; ijk) = \delta \quad (1)$$

is finite, the (infinite) graph is said to have negative or hyperbolic curvature. Other definitions of curvature count the triangles (or other polygons) that meet at each vertex[20], but these define a local, not global, curvature and can be argued to be unrelated to the global performance of networks.

For a finite graph, Eq.(1) is trivially finite and the Gromov curvature has to be modified. We introduce the concept of the ‘‘curvature plot’’ of a network: for every triangle $\Delta = (ijk)$ we plot δ_Δ vs L_Δ where

$$\begin{aligned} \delta_\Delta &= \min_m D(m; ijk) \\ L_\Delta &= \min[d(ij), d(jk), d(ki)]. \end{aligned} \quad (2)$$

This yields $P_L(\delta)$, the probability distribution for δ at fixed L . If the peak of $P_L(\delta)$ is at $\delta = \delta_p(L)$, the network is flat (negatively curved) if $\delta_p(L)$ increases linearly (sublinearly) with L [21]. Since we use the peak of the distribution instead of the maximum as in Eq.(1), statistical sampling of triangles is sufficient.

Figure 2 shows the curvature plot for network 7018(AT&T) from the Rocketfuel database [13] (see Figure 1). The metric used is the ‘hop metric’, where each edge of the graph has unit length. This is a common metric that best illustrates the geometrical properties of the graph, including the ‘small world’ property [22]. The networks in this database are at the IP layer and describe the IP port to IP port connectivity of the network. A sharp ridge is seen along the curve $\delta_p(L)$. The ridge is a straight line through the origin for the triangular lattice but bends over parallel to the L -axis for the 7018 network ($P_L(\delta)$ is zero for $\delta > 3$ for all L , though the diameter of the network is 12). For all the networks in the database, we have verified that the measured δ ’s do not exceed 3, even though the network diameters range from 12 to 14 (with the exception of 4755/VSNL whose diameter is 6, but whose ratio diameter/ δ is even bigger, 6). The ratio of 3/12 or 25% is comfortably within the theoretical bound for scaled hyperbolic graphs[24].

As another manifestation of the curvature, Figure 3 shows the average δ for each L , $E[\delta](L)$, for all ten networks in the Rocketfuel database. The plots saturate for relatively small L . The figure also shows $E[\delta](L)$ for the Barabasi-Albert model [4] and a Watts-Strogatz type model [3]; although both of these models exhibit small world behavior, we see that only the first has negative curvature as defined in Eq.(1). The plot for the Watts-Strogatz graph shows signs of saturation for large L , but the size of this graph was chosen so that it was already well in the small world regime [3, 9].

Turning to the performance implications of hyperbolic curvature, the simplest graphs with (constant) negative curvature are the hyperbolic grids $X_{p,q}$ consisting of q regular p -gons at each vertex when $(p-2)(q-2) > 4$. [25]

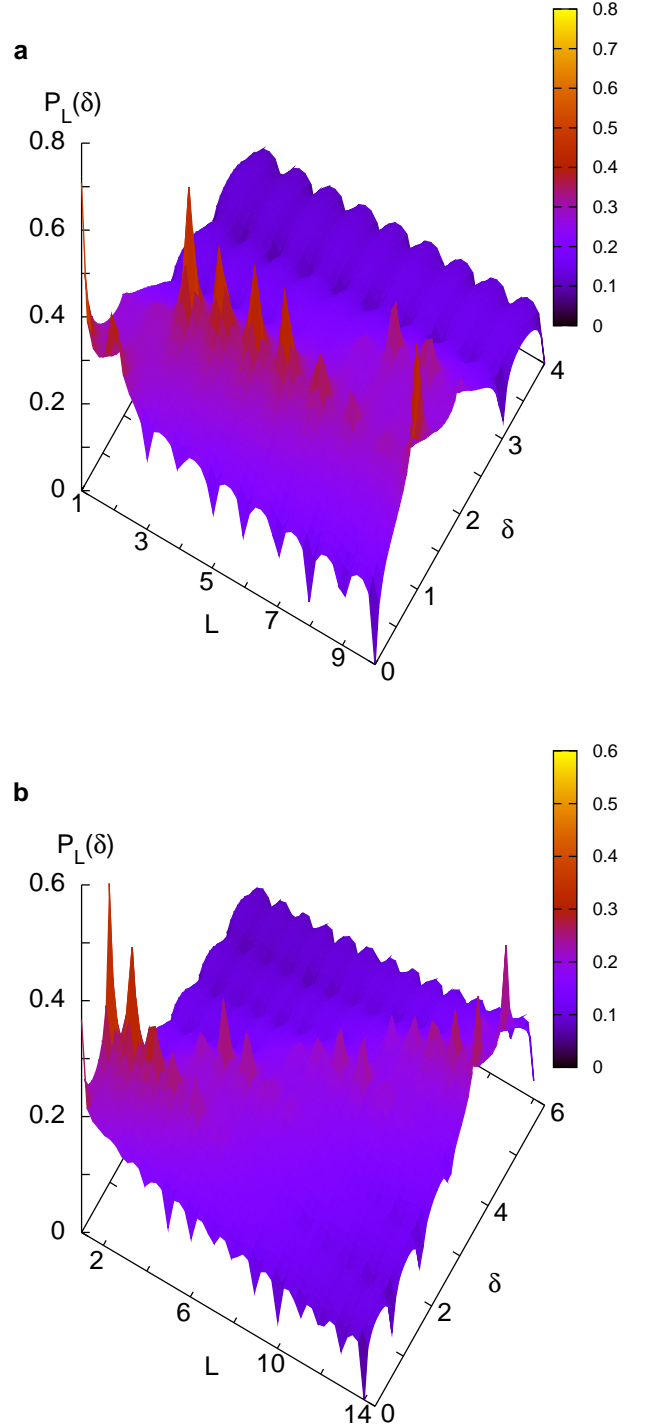


FIG. 2: (a) Probability $P_L(\delta)$ for randomly chosen triangles whose shortest side is L to have a given δ as defined in Eq.(2) for the network 7018(AT&T network) which has 10152 nodes and 14319 links and diameter 12. The quantities δ and L are restricted to integers, and the smooth plot is by interpolation. (b) Similar to (a), for a (flat) triangular lattice with 469 nodes and 1260 links. (The smaller number of nodes is sufficient for comparing with (a) since the range for L is large due to the absence of the small world effect.)




FIG. 3: The average δ as a function of L , $E[\delta](L)$, for the 10 IP-layer networks studied here, and for the Barabasi-Albert model with $k = 2$ and $N = 10000$ (11th curve) and the hyperbolic grid $X_{3,7}$ (12th curve). On the other hand, a Watts-Strogatz type model on a square lattice with $N = 6400$, open boundary conditions and 5% extra random connections (13th curve) and two flat grids (the triangular lattice with diameter 29 and the square lattice with diameter 154) are also shown.

(When $(p-2)(q-2) = 4$, the graph is flat.) We construct finite hyperbolic grids by truncating to n hops from the center. The number of nodes N in the graph increases exponentially as n is increased. With unit demand between all node pairs and the traffic between two nodes traveling along a geodesic connecting them (evenly distributed over all geodesics in case of ties), we have verified numerically that the load at the center of the graph scales with the number of nodes N in the graph as

$$L_c(N) \sim N^2. \quad (3)$$

The same result can be obtained analytically for the continuum Poincaré disk truncated to a radius $r < 1$, converted to a graph by introducing a uniform distribution of nodes with each node connected to its neighbors.[27] By contrast, it is not hard to verify that $L_c(N) \sim N^{1.5}$ for a Euclidean graph. Physically, this is because the traffic from the $\sim N$ nodes on the left of a Euclidean lattice to the $\sim N$ nodes on the right flows through the center across a line of length $\sim \sqrt{N}$, whereas for a hyperbolic graph it is pulled inwards and flows within an $O(1)$ distance from the center. Figure 4 shows the load at the node with the highest load versus N for all the networks in the Rocketfuel database, demonstrating $\sim N^2$ scaling. The figure also shows results for the Barabasi-Albert and Watts-Strogatz models; we see that the first shows $\sim N^2$ scaling but the second does not, confirming our earlier conclusion that the latter is a poorer fit to Internet-type large-scale networks.

There are two points worth noting. First, one might wonder whether the concentration of geodesics and load




FIG. 4: Plot of the maximum load $L_c(N)$ for each network in the Rocketfuel database as a function of the number of nodes N in the network. Also shown are the maximum load for the hyperbolic grid $X_{3,7}$, the Barabasi-Albert model with $k = 2$, the Watts-Strogatz model and a triangular lattice, for various N . The dashed lines have slopes of 2.0 and 1.5, corresponding to the hyperbolic and Euclidean cases respectively.




FIG. 5: **Taxonomy:** Taxonomy of key characteristics of networks and their overlaps in a schematic diagram. “Hairy” as used in this figure, refers to the simple mechanism of making a grid power-law by adding to each node a set of singly-connected nodes (hairs) whose number is drawn from any desired power-law distribution. PLDD refers to power law degree distributions.

near the center is trivial because the networks we have studied are almost simple trees. However, the ratio of the number of edges to nodes in these networks ranges from 1.27 to 2.72, showing that they are far from being trees. Second, as the example of hyperbolic grids demonstrate, one can construct graphs where every node

has the same degree, but which exhibit the ‘small world’ property *and* show N^2 scaling of load. Thus although the networks we have studied do seem to have power-law degree distributions, hyperbolicity is a nontrivial and general property that is distinct from their degree distribution and — based on the $\sim N^2$ scaling of the previous paragraph — can significantly impact performance. Figure 5 summarizes the relationship between several key characteristics discussed in the literature in the context of large-scale complex networks. We observe that hyperbolicity entails small world behavior, a fundamental property of networks.

Our results suggest that, counter-balancing the positive benefits of hyperbolicity such as the small world property, core congestion is a structural *problem* due such

hyperbolicity that grows more acute as the network grows in size. As long as routing protocols use geodesics in one form or another, whether in intra-domain, inter-domain or other forms of routing, congestion is a natural consequence of this intrinsic structural feature of networks. Using ‘ $(1 + \epsilon)$ routing’, in which traffic between nodes is *not* routed along the geodesic(s) between them but is deliberately sent on slightly longer paths, would in fact alleviate core congestion. This is a phenomenon familiar from vehicular traffic: the shortest routes using expressways can become so overcrowded that indirect and longer paths through backroads become faster.

This work was funded by AFOSR grant FA9550-08-1-0064.

-
- [1] This is the same as the shortest path length when each edge is given weight 1.
 - [2] P. Erdos, A. Renyi, *Publicationes Mathematicae* **6**, 290 (1959).
 - [3] D.J. Watts, S.H. Strogatz, *Nature* **393**, 409 (1998).
 - [4] R. Albert, H. Jeong, A-L. Barabasi, *Nature* **401**, 130 (1999).
 - [5] P.L. Krapivsky, S. Redner, F. Leyvraz, *Phys. Rev. Lett.* **85**, 4629 (2000).
 - [6] S.N. Dorogovtsev, J.F.F. Mendes, A. N. Samukhin, *Phys. Rev. E* **64**, 066110 (2001).
 - [7] M. Faloutsos, P. Faloutsos, C. Faloutsos, *Proc. ACM SIGCOMM’99*, 251 (1999).
 - [8] A-L. Barabasi, H. Jeong, Z. Neda, E. Ravasz, A. Schubert, T. Vicsek, *Physica A* **311**, 590 (2002).
 - [9] M. E. J. Newman, *SIAM Review* **45**, 167 (2003).
 - [10] J. G. White, E. Southgate, J.N. Thomson, S. Brenner, *Phil. Trans. Roy. Soc. Lond.* **B 314**, 1 (1986).
 - [11] In this context, the ‘network core’ is the set of nodes that are at the intersection of the majority of geodesics (shortest paths), while the ‘network edge’ is the set where this property is minimal.
 - [12] Y. Baryshnikov, *Workshop Stochastic Geom. Teletraffic*, Eindhoven (2002).
 - [13] N. Spring, R. Mahajan, D. Wetherall, *Proc. ACM/SIGCOMM ’02*, 133 (2002). Data archive at <http://www.cs.washington.edu/research/networking/rocketfuel/>.
 - [14] D. Krioukov, F. Papadopoulos, M. Boguna, A. Vahdat, <http://arxiv.org/abs/0805.1266> (2008).
 - [15] M. Gromov, *Essays in Group Theory*, Math. Sci. Res. Inst. Publ. **8**, 75 (Springer, New York, 1987).
 - [16] K.-I. Goh, B. Kahng, and D. Kim, *Phys. Rev. Lett.* **87**, 278701 (2001).
 - [17] K.-I. Goh, H. Jeong, B. Kahng, and D. Kim, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 12583 (2002).
 - [18] M. Barthelemy, *Phys. Rev. Lett.* **91**, 189803 (2003).
 - [19] O. Narayan and I. Saniee, <http://arxiv.org/abs/0906.4138> (2009).
 - [20] For example, the Descartes or Gauss-Bonnet curvature defined in terms of the angular defect or excess from 2π as measured at each node.
 - [21] By choosing m to be on the sides of the triangle, it is easy to verify that $\delta_\Delta \leq L_\Delta/2$, so that positive curvature cannot be seen with this test or Eq.(1).
 - [22] While routing in the Internet can involve link metrics other than the hop metric, and careful sizing of links and assignment of such metrics can avoid congestion in pre-specified hot spots in the network, such measures often create hot spots elsewhere.[23].
 - [23] Z. Wang, J. Crowcroft, *ACM SIGCOMM Computer Comm. Rev.* **22**, 63 (1992).
 - [24] E. Jonckheere, P. Lohsoonthorn, F. Bonahon, *J. Graph Th.* **57**, 157 (2007).
 - [25] These grids correspond to Cayley graphs of the fundamental groups of surfaces (or graphs) with genus g when $p = q = 4g$. [26]
 - [26] H.S.M. Coxeter, *Proc. Intern. Cong. Math.* 1954 v. 3, 155 (North Holland, Amsterdam, 1956).
 - [27] This model is applicable to a hyperbolic network if the spacing between nodes is small compared to its radius of curvature.